

Congestion-Dependent Pricing of Network Services

Ioannis Ch. Paschalidis, *Member, IEEE*, and John N. Tsitsiklis, *Fellow, IEEE*

Abstract—We consider a service provider (SP) who provides access to a communication network or some other form of on-line services. Users initiate calls that belong to a set of diverse service classes, differing in resource requirements, demand pattern, and call duration. The SP charges a fee per call, which can depend on the current congestion level, and which affects users' demand for calls. We provide a dynamic programming formulation of the problems of revenue and welfare maximization, and derive some qualitative properties of the optimal solution. We also provide a number of approximate approaches, together with an analysis that indicates that near-optimality is obtained for the case of many, relatively small, users. In particular, we show analytically as well as computationally, that the performance of an optimal pricing strategy is closely matched by a suitably chosen static price, which does not depend on instantaneous congestion. This indicates that the easily implementable *time-of-day* pricing will often suffice. Throughout, we compare the alternative formulations involving revenue or welfare maximization, respectively, and draw some qualitative conclusions.

Index Terms—Dynamic programming, internet economics, loss networks, revenue management.

I. INTRODUCTION

IN THIS paper, we consider a service provider (SP) who provides access to a communication network or some other form of on-line services. Users access the network and initiate calls that belong to a set of diverse service classes, differing in resource requirements, demand pattern, and call duration. The SP charges a fee per call, which can depend on the current congestion level, and which affects users' demand for calls. We are interested in the problem of determining an optimal pricing strategy, that maximizes an appropriate performance measure such as social welfare or the provider's revenue. As discussed in the sequel (Section III), our model applies to a variety of situations in which a user population shares dynamically a limited resource, but our primary motivation comes from the context of communication network services, e.g., provided through the Internet.

For the case of a revenue maximizing provider, we have a problem of *yield management*, similar to the problems that arise in service industries (e.g., airlines [1]). The common element in

such problems is that the marginal cost of serving an additional customer, e.g., an airline passenger or a new call, is negligible once a flight has been scheduled or a communications infrastructure is in place.

The pricing problem in communication networks, such as the Internet, is complex and multifaceted, because one needs to take into account engineering issues (compatibility with existing or planned protocols), financial issues (cost recovery), a diverse set of desired services (e.g., elastic versus real-time traffic), and other considerations (such as simplicity, social welfare, etc.). The current Internet relies on technical means to prevent congestion (the TCP protocol), but includes no mechanisms for ensuring quality of service (QoS) guarantees or for delivering service to those users who need it most. Pricing mechanisms can overcome these shortcomings, resulting in more efficient resource allocation, by charging users on the basis of the congestion that they cause.

In this spirit, several models have been put forth. MacKie-Mason and Varian [2] have proposed a “smart market” where packets bid for transport while the network only serves packets with bids above a certain cutoff amount, depending on the level of congestion. Kelly *et al.* [3], [4] consider charges that increase with either realized flow rate or with the “share” of the network consumed by a traffic stream. Gibbens and Kelly [5] describe yet another scheme for packet-based pricing as an incentive for more efficient flow control. Clark [6] proposes an expected capacity-based pricing scheme where users are charged ahead of time on the basis of the expectation that they have of network usage and excess packets are dropped at times of congestion.

The emergence of real-time traffic substantially complicates the picture and requires QoS measures much harder to analyze (see, e.g., Kelly [7], Bertsimas *et al.* [8], [9], and Paschalidis [10]). More importantly, users who access the network wish to complete certain tasks (e.g., send an e-mail, access a Web page, place an “internet phone” call) and packets are completely transparent to them. Thus, it is not clear whether packet-based pricing schemes are always appropriate. Kelly [11] and Courcoubetis *et al.* [12] propose the pricing of real-time traffic with QoS requirements, in terms of its “effective bandwidth,” and provide approximations that only involve time and volume charges.

A central question that has received little formal treatment relates to determining the right time scale over which prices should evolve. For example, prices could stay fixed, or they might depend on slowly-varying parameters that capture the prevailing operating conditions (as in time-of-day pricing), or they could respond to statistical fluctuations of congestion. This is one of the main issues addressed in this paper. We will establish that for both revenue and welfare maximization and under stationary demand, fixed (static) prices are asymptotically optimal in a number of limiting regimes including light traffic,

Manuscript received November 10, 1998; revised August 4, 1999; approved by IEEE/ACM TRANSACTIONS ON NETWORKING Editor S. Low. The work of I. Ch. Paschalidis was supported in part by the Internet & Telecoms Convergence Consortium, by the National Science Foundation under Grant NCR-9706148 and Grant ACL-9873339, and by the Nokia Research Center, Boston, MA. The work of J. N. Tsitsiklis was supported in part by the National Science Foundation under Contract ACI-9873339 and under Grant INT-9725223.

I. Ch. Paschalidis is with the Department of Manufacturing Engineering, Boston University, Boston, MA 02215 USA (e-mail: yannisp@bu.edu).

J. N. Tsitsiklis is with the Laboratory for Information and Decision Systems, Massachusetts Institute of Technology, Cambridge, MA 02139 USA (e-mail: jnt@mit.edu).

Publisher Item Identifier S 1063-6692(00)03319-7.

heavy traffic, and a regime of many, relatively small, users, which is of more practical importance in a network environment such as the Internet. This indicates that when demand statistics are slowly varying *time-of-day* pricing will often suffice.

On the technical side, the problem that we study is structurally similar to arrival and service rate control problems for single server queueing systems; see, e.g., Lippman and Stidham [13] as well as the work of Subramanian *et al.* [14] on yield management of airline reservations. It also has similarities with the network pricing problem in Gallego and Van Ryzin [15]. Their paper considers a finite horizon formulation (versus our infinite horizon average-cost setup). Our problem is also related to problems of admission control in loss networks (see, e.g., Key [16], Ott and Krishnan [17], and Ross [18]). It is different, however, because this literature assumes that the prices are fixed and is only concerned with admission decisions, while we wish to study optimal or near-optimal pricing schemes. Another difference with most of the literature is that we use a decision-theoretic framework under an explicit model of users' reaction to prices (demand functions). Similar demand functions have been used in Low and Varaiya [19] under a somewhat different model.

The remainder of the paper is organized as follows. In Section II, we formulate the two problems of revenue and welfare maximization. In Section III, we discuss the applicability of our model. In Section IV, we indicate how to obtain an optimal dynamic pricing policy using dynamic programming and derive some qualitative properties of the optimal solution. After a brief discussion of static (congestion-independent) pricing, in Section V, we proceed to develop bounds and approximations. In Section VI, we derive an easily computable upper bound on the optimal performance. Then, using this bound, in Section VII we establish asymptotic optimality of the static policy in light traffic, heavy traffic, and a regime involving many and relatively small users. We also draw a number of qualitative conclusions. In Sections VIII and IX, we discuss briefly a number of methods for obtaining approximately optimal static and dynamic pricing policies, respectively. In Section X, we present numerical results, including a comparison of different approaches. Moreover, we compute approximately optimal dynamic policies for some large scale examples. In Section XI, we discuss the case of slowly-varying or imprecisely-known demand statistics. Conclusions and extensions are in Section XII.

II. PROBLEM FORMULATION

In this section, we introduce a model for the operation of an SP. We assume that the SP has a total amount R of some resource and that each service request ("call") needs a certain amount of that resource. We will be referring to this resource as "bandwidth" and to R as the "capacity" although, as will be discussed later, other interpretations are also possible.

We assume that calls belong to M different classes. Calls of class $i = 1, \dots, M$, arrive according to a Poisson process and stay connected for a time interval which is exponentially distributed with rate μ_i . Let $\boldsymbol{\mu} = (\mu_1, \dots, \mu_M)$. Upon arrival, a call of class i pays a fee of u_i ; we denote $\mathbf{u} = (u_1, \dots, u_M)$. We assume that there is a known *demand function* $\lambda_i(u_i)$, which determines the arrival rate of class i calls, as a function of the price

u_i . The following assumption will remain in effect throughout the paper.

Assumption A: For every i , there exists a price $u_{i,\max}$ beyond which the demand $\lambda_i(u_i)$ becomes zero. Furthermore, the function $\lambda_i(u_i)$ is continuous and strictly decreasing in the range $u_i \in [0, u_{i,\max}]$.

We will write $\boldsymbol{\lambda}(\mathbf{u}) = (\lambda_1(u_1), \dots, \lambda_M(u_M))$, and we will be denoting by $\boldsymbol{\lambda}_0$ the vector of arrival rates when the price is zero, i.e., $\boldsymbol{\lambda}_0 = (\lambda_{0,1}, \dots, \lambda_{0,M}) \triangleq \boldsymbol{\lambda}(\mathbf{0})$. By monotonicity, $\boldsymbol{\lambda}_0 \geq \boldsymbol{\lambda}(\mathbf{u})$ for all nonnegative price vectors \mathbf{u} .

Let $n_i(t)$ be the number of class i calls that are in progress at time t . Since $n_i(t)$ is discontinuous at the times of call arrivals and departures, we adopt the convention that $n_i(t)$ is a right-continuous function of t . We will be writing $\mathbf{N}(t) = (n_1(t), \dots, n_M(t))$. An incoming class i call requires r_i units of bandwidth and is only accepted if that bandwidth is available, that is, if $\mathbf{N}(t)' \mathbf{r} + r_i \leq R$, where $\mathbf{r} = (r_1, \dots, r_M)$ and prime denotes transpose. A rejected call gets a busy signal and is lost for the system.

A pricing policy is a rule that determines the current price vector $\mathbf{u} = (u_1, \dots, u_M)$ as a function of the current state $\mathbf{N}(t)$. We therefore use the notation $\mathbf{u}(t)$ to indicate that the price vector is time-dependent. Without loss of generality, we can assume that whenever $\mathbf{N}(t)' \mathbf{r} + r_i > R$, we set $u_i = u_{i,\max}$ and there are no class i arrivals.

A. Formulation of the Revenue Maximization Problem

Under any given pricing policy, the system evolves as a continuous-time Markov chain with state $\mathbf{N}(t)$. Given the current time t and price $\mathbf{u}(t)$, and for small δ , there is a probability approximately equal to $\lambda_i(u_i(t))\delta$ that there is a class i arrival during the next δ time units, and the expected revenue from class i arrivals during that interval is approximately $\delta u_i(t) \lambda_i(u_i(t))$. Thus, the expected long-term average revenue is given by

$$\lim_{T \rightarrow \infty} \frac{1}{T} \mathbf{E} \left[\int_0^T \boldsymbol{\lambda}(\mathbf{u}(t))' \mathbf{u}(t) dt \right]. \quad (1)$$

(The above limit is easily seen to exist for any pricing policy, because the state $\mathbf{N} = \mathbf{0}$, corresponding to an empty system, is recurrent.) We are interested in identifying a pricing policy that maximizes the above quantity.

B. Formulation of the Welfare Maximization Problem

It is also of interest to consider the case where social welfare is maximized. Toward this purpose, we need more specific assumptions on the nature of user demand, that will allow us to make inferences on user utility or welfare.

We interpret our demand model as follows. *Potential* calls of class i are born according to a Poisson process with constant rate $\lambda_{0,i}$. (This is the maximum arrival rate introduced earlier.) A potential call of class i , if it goes through, results in a user utility of U_i , where U_i is a nonnegative random variable taking values in the range $[0, u_{i,\max}]$, and which is described by a continuous probability density function $f_i(u_i)$. We assume that a potential call will go through if and only if the utility U_i exceeds

the prevailing price u_i . This implies that class i calls are realized according to a randomly modulated Poisson process with rate $\lambda_i(u_i(t)) = \lambda_{0,i} \mathbf{P}[U_i \geq u_i(t)]$. Furthermore, the expected utility conditioned on the fact that a call has been established, under a current price of u_i , is equal to $\mathbf{E}[U_i | U_i \geq u_i]$. We conclude that the expected long-term average rate at which utility is generated is given by

$$\lim_{T \rightarrow \infty} \frac{1}{T} \sum_{i=1}^M \mathbf{E} \left[\int_0^T \lambda_i(u_i(t)) \mathbf{E}[U_i | U_i \geq u_i(t)] dt \right]. \quad (2)$$

This objective is of the same form as in revenue maximization, except that the instantaneous reward rate $\lambda_i(u_i)u_i$ of class i is replaced by $\lambda_i(u_i) \mathbf{E}[U_i | U_i \geq u_i]$. Thus, the two problems can be approached using the same set of tools.

The following two formulas are immediate consequences of the above description and will be used later on:

$$\lambda_i(u_i) = \lambda_{0,i} \int_{u_i}^{u_{i,\max}} f_i(v) dv \quad (3)$$

$$\lambda_i(u_i) \mathbf{E}[U_i | U_i \geq u_i] = \lambda_{0,i} \int_{u_i}^{u_{i,\max}} f_i(v) v dv. \quad (4)$$

Example: Suppose that the utility derived from a call of class i is uniformly distributed in the range $[0, u_{i,\max}]$. Then, it is easily checked that

$$\lambda_i(u_i) = \lambda_{0,i} \left(1 - \frac{u_i}{u_{i,\max}} \right) \quad (5)$$

$$\mathbf{E}[U_i | U_i \geq u_i] = \frac{u_i + u_{i,\max}}{2}. \quad (6)$$

We observe that the reward rate for class i calls is a concave quadratic function of the prevailing price u_i for either case of revenue or welfare maximization.

Remark: Due to Assumption A the function $\lambda_i(u_i)$ has an inverse $u_i(\lambda_i)$, defined on $[0, \lambda_{0,i}]$. Furthermore, $u_i(\lambda_i)$ is continuous and strictly decreasing. This is a usual assumption in the yield management literature (see Gallego and Van Ryzin [15]) and allows us to view the arrival rates λ_i as the SP's decision variables. Sometimes, we will also assume that the instantaneous reward rate $\lambda_i u_i(\lambda_i)$ (in the revenue maximization problem) or $\lambda_i \mathbf{E}[U_i | U_i \geq u_i(\lambda_i)]$ (in the welfare maximization problem) is concave. (This is always the case with the linear demand functions in the example above.) With these assumptions, the demand function is *regular*, in the terminology of [15].

III. APPLICABILITY OF THE MODEL

Our model is relevant to a variety of contexts. The following is a partial list of possibilities.

- 1) A network access provider has a finite modem pool and serves a large user population. The price charged for a call can depend on the present level of utilization in the system. In this example, all calls can be viewed as belonging to a single class, since they all consume one unit of the available resource (one modem).
- 2) A network provider provides a menu of possible connection types, reflecting the nature of the session (e.g., voice versus video) and the desired QoS. The assumption is

that the requirements of any particular service type can be characterized by a single number (e.g., some form of effective bandwidth) and that the resources available to the provider can be similarly described in terms of total available effective bandwidth.

- 3) A server provides content (e.g., data base access, on-line financial information), or access to a computer program ("applications on tap") possibly through the Internet. Requests for content ("calls") can be of different types in terms of expected duration and in terms of how much of the server capacity needs to be reserved in order to guarantee the desired QoS.

Our framework incorporates a number of assumptions. The most important one is that the number of calls that can be accommodated obeys a *linear* constraint. This is natural if the constraining resource is a simple quantity like bandwidth. More generally, any system can be described by its "admission region," defined as the set of all vectors (n_1, \dots, n_M) such that the system is able to simultaneously accommodate n_i calls of each class i , at an acceptable QoS. This admission region can be quite difficult to evaluate, especially if there is a complex interaction between different calls through statistical multiplexing (see [7], [10]), and may depend on low-level details such as the scheduling algorithms used at multiplexers. For the purpose of economic analysis, an exact but complex description of the admission region is unlikely to be useful. Instead, it is reasonable to aim at a single scalar that summarizes the resource requirements of any given class, and impose a linear constraint, as in our model. This is the approach advocated by Kelly and coworkers [7], based on a linearization of the constraints defining the admission region in the vicinity of a likely operating point.

We have also assumed that the demand functions $\lambda_i(u_i)$ are known. There are some current research activities [20] that aim at understanding the dependence of demand on price, but it is plausible that the demand curves assumed here will never be known with any great precision. Nevertheless, as in much of economics research, one can postulate the existence of such demand curves and derive valuable insights through mathematical analysis. These insights can then lead to "adaptive" methods that rely on observed behavior rather than on given demand models.

A limitation of our model is that it does not incorporate any substitution effects either between classes or in time; see Courcoubetis and Reiman [21] for a model that does.

We finally note that our model involves a single charge at the time that a call is admitted. Even though certain types of "unlimited connect time" services are commercially available, this is not economically sensible, because it does not provide an incentive for the users to stop consuming network resources. At the opposite extreme, one might consider connection time charges at a rate that is continuously updated on the basis of the present state of the system. This could be undesirable to users because the total charge for a call would not be known at the time of the connection. As a compromise, the connect time charge rate for a class i user, denoted by c_i , could be determined at the time that the call is established, and then remain locked throughout the call. In that case, the expected user payment becomes $u_i + (c_i/\mu_i)$. Mathematically, this is equivalent

to having the user make the full payment up front, and no separate analysis is needed. More complex models, not analyzed in this paper, are possible, e.g., when the mean duration $1/\mu_i$ depends on the connect time charge rate c_i .

IV. OPTIMAL DYNAMIC POLICIES

In this section, we show how to obtain an optimal pricing policy using *dynamic programming (DP)*, and we derive some of its properties. To keep the exposition simple, we mainly discuss the maximization of the provider's revenue. At the end of the section, we comment on the corresponding results for the welfare maximization formulation.

A. Dynamic Programming Formulation

The problem introduced in Section II is a finite-state, continuous-time, average reward DP problem. Note that the set $\mathcal{U} = \{\mathbf{u} \mid 0 \leq u_i \leq u_{i,\max} \forall i\}$ of possible price vectors is compact and that all states communicate, i.e., for each pair of states \mathbf{N}, \mathbf{N}' , there exists a policy under which we can eventually reach \mathbf{N}' starting from \mathbf{N} . Having assumed that the demand functions $\lambda_i(u_i)$ are continuous, the transition rates as well as the reward rate $\sum_i \lambda_i(u_i)u_i$ are continuous in the decision variables. Moreover, the reward rate and the expected holding time at each state \mathbf{N} are bounded functions of \mathbf{u} . The same is true for the total transition rate out of any state. Under these assumptions, the standard DP theory applies (see [22] and [23]) and asserts that there exists a stationary policy which is optimal.

The process $\mathbf{N}(t)$ is a continuous-time Markov chain. Since the total transition rate out of any state is bounded by $\nu = \sum_{i=1}^M (\lambda_{0,i} + \mu_i [R/r_i])$, this Markov chain can be uniformized, leading to a Bellman equation of the form

$$\begin{aligned} J^* + h(\mathbf{N}) = & \max_{\mathbf{u} \in \mathcal{U}} \left[\sum_{i \notin C(\mathbf{N})} \lambda_i(u_i)u_i \right. \\ & + \sum_{i \notin C(\mathbf{N})} \frac{\lambda_i(u_i)}{\nu} h(\mathbf{N} + \mathbf{e}_i) \\ & + \sum_{i=1}^M \frac{n_i \mu_i}{\nu} h(\mathbf{N} - \mathbf{e}_i) \\ & \left. + \left(1 - \sum_{i \notin C(\mathbf{N})} \frac{\lambda_i(u_i)}{\nu} - \sum_{i=1}^M \frac{n_i \mu_i}{\nu} \right) h(\mathbf{N}) \right]. \end{aligned} \quad (7)$$

Here, $C(\mathbf{N}) = \{i \mid (\mathbf{N} + \mathbf{e}_i)' \mathbf{r} > R\}$ is the set of classes whose calls cannot be admitted in state \mathbf{N} . We impose the condition $h(\mathbf{0}) = 0$, in which case Bellman's equation has a unique solution [in the unknowns J^* and $h(\cdot)$]. Once Bellman's equation is solved, an optimal policy is readily obtained by choosing at each state \mathbf{N} a price vector \mathbf{u} that maximizes the right-hand side in (7). The solution to Bellman's equation has the following interpretation: the scalar J^* is the optimal expected revenue per unit time, and $h(\mathbf{N})$ is the *relative reward* in state \mathbf{N} . In particular, consider an optimal policy that attains the maximum in (7) for every state \mathbf{N} . If we follow this policy starting from state \mathbf{N}' or

state \mathbf{N} , the expectation of the *difference* in total rewards (over the infinite horizon) is equal to $(h(\mathbf{N}') - h(\mathbf{N}))/\nu$.

The solution to Bellman's equation and a resulting optimal policy can be computed using classical DP algorithms. However, the computational complexity increases with the size of the state space, which is exponential in the number of classes M . For this reason, an exact solution using DP is feasible only when the number of classes is quite small.

B. Some General Properties

Our first result establishes the monotonicity of the relative rewards. It corresponds to the intuitive fact that it is always more desirable to have more free resources, as they lead to additional revenue generating opportunities in the future. The proof uses a coupling argument. In order to carry out this proof (as well as the proof of Theorem 5, later on), we need some notation and a few facts from DP theory.

We define the DP operator T , which maps the set of functions on the state space into itself, as follows: for any such function f , $(Tf)(\mathbf{N})$ is defined to be equal to the right-hand side of (7), with h replaced by f . In particular, (7) can be written as

$$J^* + h(\mathbf{N}) = (T\mathbf{h})(\mathbf{N}). \quad (8)$$

For any policy π , we define an operator T_π similarly, except that instead of maximizing with respect to \mathbf{u} , we use the price vector determined by the policy π .

Let T^k denote the composition of k copies of T . Then, $(T^k f)(\mathbf{N})$ is equal to the optimal total expected reward in a k -stage problem with starting state \mathbf{N} , terminal reward function f , and reward per stage equal to $\sum_i \lambda_i(u_i)u_i$. (A "stage" here refers to a transition in the uniformized chain.) The composition T_π^k is defined similarly and admits the same interpretation, except that it refers to expected reward under the fixed policy π .

We will need the following facts, which are true for average reward problems with a state which is recurrent under any policy. For any function f , and every state \mathbf{N} , we have

$$\lim_{k \rightarrow \infty} [(T^k f)(\mathbf{N}) - kJ^*] = h(\mathbf{N}). \quad (9)$$

Let π be a policy which at every state \mathbf{N} attains the maximum of the right-hand side in the Bellman equation. Such a policy is optimal and, furthermore,

$$\lim_{k \rightarrow \infty} [(T_\pi^k f)(\mathbf{N}) - kJ^*] = h(\mathbf{N}) \quad (10)$$

for every function f and state \mathbf{N} .

Theorem 1 [Monotonicity of $h(\mathbf{N})$]: For all j and all \mathbf{N} such that $(\mathbf{N} + \mathbf{e}_j)' \mathbf{r} \leq R$, we have $h(\mathbf{N}) \geq h(\mathbf{N} + \mathbf{e}_j)$, where \mathbf{e}_j denotes the j th unit vector.

Proof: Let π be an optimal policy that attains the maximum in the right-hand side of Bellman's equation for every state \mathbf{N} . Consider a k -stage (k transitions of the uniformized chain) problem with terminal reward function $\mathbf{0}$, identically equal to zero. We consider the system starting from two different initial states. The first, which we refer to as System A, starts from state $\mathbf{N} + \mathbf{e}_j$ and follows the optimal policy π . Its k -stage expected reward is equal to $(T_\pi^k \mathbf{0})(\mathbf{N} + \mathbf{e}_j)$. The second, which we refer to as System B, starts from state \mathbf{N} , but at any time uses the

same prices as System A. The statistics of the arrival processes for the two systems are the same, and by defining the two processes on a common probability space, we can assume that the actual arrivals are the same. We can also identify the common \mathbf{N} customers in the two systems and assume that they depart at the exact same instances. Then, the state of System B is at all times less than or equal to the state of System A, and in particular, whenever a customer is admitted in System A, a customer of the same class is also admitted to System B, at the same price. Thus, the revenue in System A is equal to the revenue in System B, and the same remains true after we take expectations. But the expected optimal k -stage revenue for System B is at least as large and this shows that $(T^k \mathbf{0})(\mathbf{N}) - (T^k \mathbf{0})(\mathbf{N} + \mathbf{e}_j) \geq 0$. We take the limit as $k \rightarrow \infty$ and use (9) and (10) to conclude that $h(\mathbf{N}) - h(\mathbf{N} + \mathbf{e}_j) \geq 0$. ■

We now record some observations for the case where the SP has infinite capacity.

Theorem 2 (The Infinite Bandwidth Case): If there are no capacity constraints ($R = \infty$), the optimal revenue is given by

$$J_\infty = \max_{\mathbf{u} \in \mathcal{U}} \sum_{i=1}^M \lambda_i(u_i) u_i$$

and the optimal price vector is some constant \mathbf{u}_∞ that does not depend on the state \mathbf{N} . Furthermore, we have $J^* \leq J_\infty$.

We now show that resource limitations always result in higher prices in comparison to the unconstrained case.

Theorem 3: There exists an optimal policy \mathbf{u}^* such that for every state \mathbf{N} , we have $\mathbf{u}^*(\mathbf{N}) \geq \mathbf{u}_\infty$.

Proof: Fix some state \mathbf{N} . From the Bellman equation, we see that for all $i \notin C(\mathbf{N})$, an optimal price $u_i^*(\mathbf{N})$ can be chosen by maximizing the expression

$$\lambda_i(u_i) u_i + \frac{\lambda_i(u_i)}{\nu} (h(\mathbf{N} + \mathbf{e}_i) - h(\mathbf{N})).$$

Consider a value of u_i which is less than the i th component $u_{\infty,i}$ of \mathbf{u}_∞ . Then, $\lambda_i(u_i) u_i \leq \lambda_i(u_{\infty,i}) u_{\infty,i}$, by the definition of $u_{\infty,i}$. By Theorem 1, we have $h(\mathbf{N} + \mathbf{e}_i) - h(\mathbf{N}) \leq 0$. Also, by monotonicity of the demand function, we have $\lambda_i(u_i) \geq \lambda_i(u_{\infty,i})$. Using all of the above inequalities

$$\begin{aligned} & \lambda_i(u_i) u_i + \frac{\lambda_i(u_i)}{\nu} (h(\mathbf{N} + \mathbf{e}_i) - h(\mathbf{N})) \\ & \leq \lambda_i(u_{\infty,i}) u_{\infty,i} + \frac{\lambda_i(u_{\infty,i})}{\nu} (h(\mathbf{N} + \mathbf{e}_i) - h(\mathbf{N})). \end{aligned}$$

This implies that u_i cannot be strictly better than $u_{\infty,i}$. ■

C. Price Monotonicity in the Single Class Case

We now assume that there is a single service class and that $r = 1$. (Accordingly, we simplify notation, and use n to denote the state of the system.) We wish to show that the optimal price is an increasing function of the state. We first show that the relative rewards are concave.

Theorem 4 [Concavity of $h(n)$]: For all n satisfying $0 < n < R$, we have $h(n) \geq (1/2)h(n-1) + (1/2)h(n+1)$.

Proof: Fix some n with $0 < n < R$. Let $u^*(\cdot)$ be an optimal policy that attains the maximum in the right-hand side of (7) for every state. We consider two copies of the system.

The first, which we refer to as System A, starts from state $n-1$ and follows the optimal policy. The second, which we refer to as System B, starts from state n . We identify $n-1$ of the customers in System B with those in System A, so that they have identical departure times. We refer to the remaining customer in System B as the tagged customer. System B operates as follows: at any point in time, it sets the same prices as System A does, and can therefore be assumed to observe the same sequence of arrivals and the same revenue stream. This is done until the time t that the tagged customer departs, or until the first time that System B moves to state $n+1$, whichever comes first. After that, System B follows the same optimal policy $u^*(\cdot)$, as in System A.

If the departure of the tagged customer comes first (i.e., before System B reaches state $n+1$), the two systems have the same revenue stream until time t . At time t , they are found in the same state, and so their revenue streams coincide over the infinite horizon. If on the other hand System B moves to state $n+1$ before the departure of the tagged customer, which will happen with some probability $p < 1$, then Systems A and B are found in states n and $n+1$, respectively. From that time on, both systems follow the same optimal policy, so that the difference in their expected total future rewards is $(h(n+1) - h(n))/\nu$. We conclude that the difference of the expected total future reward (starting from time zero) of System B minus that of System A, is $p(h(n+1) - h(n))/\nu$. If System B was using an optimal policy at all times, the expected future revenue would be no smaller, which establishes that

$$h(n) - h(n-1) \geq p(h(n+1) - h(n)).$$

Using the inequalities $h(n+1) - h(n) \leq 0$ (cf. Theorem 1) and $p \leq 1$, we obtain

$$h(n) - h(n-1) \geq h(n+1) - h(n)$$

which is the desired inequality. ■

We next establish that in the single class case, optimal prices reflect the level of congestion, that is, the revenue maximizing SP will raise prices as the system becomes more congested. To that end, we use the result of Theorem 4.

Theorem 5 (Monotonicity of Optimal Prices): There exists an optimal policy u^* with the property $u^*(n) \geq u^*(n-1)$, for $n = 1, \dots, R-1$.

Proof: Since the mapping $u \mapsto \lambda$ is invertible, we may view λ as the controlled variable, instead of u . From the Bellman equation, we see that an optimal choice of λ , at state n , must maximize the function

$$F(\lambda) + \frac{\lambda}{\nu} (h(n+1) - h(n))$$

where $F(\lambda) = u(\lambda)\lambda$.

Fix some n with $0 < n < R$. Let $g_n = (h(n+1) - h(n))/\nu$, $g_{n-1} = (h(n) - h(n-1))/\nu$, and let λ_n, λ_{n-1} be maximizing arrival rates at states n and $n-1$, respectively. We distinguish two cases: $g_n = g_{n-1}$ and $g_n < g_{n-1}$. (The case $g_n > g_{n-1}$ is impossible due to Theorem 4.) If $g_n = g_{n-1}$, we can certainly set $\lambda_n = \lambda_{n-1}$. Consider now the case where $g_n < g_{n-1}$. Since λ_n is optimal at state n , we have

$$F(\lambda_n) + \lambda_n g_n \geq F(\lambda_{n-1}) + \lambda_{n-1} g_n.$$

Also, since λ_{n-1} is optimal at state $n-1$, we have

$$F(\lambda_{n-1}) + \lambda_{n-1}g_{n-1} \geq F(\lambda_n) + \lambda_n g_{n-1}.$$

These two inequalities imply that

$$(\lambda_n - \lambda_{n-1})(g_n - g_{n-1}) \geq 0$$

and therefore $\lambda_n \leq \lambda_{n-1}$. Since $\lambda(u)$ is a decreasing function of u , we can conclude that in both cases, an optimal price at state n can be chosen no smaller than the optimal price at state $n-1$. ■

Unfortunately, Theorem 5 does not extend to the multiclass case, as shown by a counterexample provided in the Appendix. This is similar to known counterexamples for the problem of optimal admission control (under fixed prices/rewards) for multiclass loss networks; see, e.g., Ross and Tsang [24] and Ross [18]. The reason behind the counterexample is the combinatorial nature of an underlying ‘‘packing’’ problem that arises from the different bandwidth requirements r_i of the two classes. It is not clear whether a similar counterexample is possible when all classes have the same bandwidth requirements. Furthermore, this combinatorial aspect should become insignificant in the limit where the capacity is large in comparison to the maximum (over all classes) bandwidth requirement.

D. Computational Example

To illustrate the structure of an optimal dynamic pricing policy, we consider an example involving a single class with a linear demand function

$$\lambda(u) = \lambda_0 - \lambda_1 u, \quad u \in [0, \lambda_0/\lambda_1]. \quad (11)$$

For a particular choice of problem data, the optimal dynamic policy is depicted in Fig. 1. As predicted by Theorem 5, it is optimal to raise the price when the system is congested, so that users are discouraged from connecting, and to lower the price to attract calls when the system is underutilized. If we were to set $R = \infty$, the optimal price u_∞ is equal to 6. We observe that this is lower than the optimal dynamic price at every state, as predicted by Theorem 2.

E. The Case of Welfare Maximization

For the case of welfare maximization, Bellman's equation remains the same, except that the reward rate $\lambda(\mathbf{u})' \mathbf{u}$ is replaced by $\sum_i \lambda_i(u_i) \mathbf{E}[U_i | U_i \geq \mathbf{u}_i]$. As in Theorem 1, the relative rewards $h(\mathbf{N})$ are again monotonically nonincreasing in \mathbf{N} , because the same proof applies. If the bandwidth is infinite, welfare is maximized by admitting every user, and the optimal price \mathbf{u}_∞ is equal to zero. When R is finite, the optimal prices are nonnegative, which provides a trivial extension of Theorem 3. For the single class case, the proof of Theorem 4 remains valid. The relative rewards are again concave in the state variable n , which results in monotonically nondecreasing optimal prices, as in Theorem 5.

F. Multiple Classes with Identical Characteristics

Suppose that all classes have identical technical characteristics, namely, $\mu_i = \mu$ and $r_i = r = 1$, for all i . The classes can




Fig. 1. Optimal dynamic pricing policy when $R = 30$, $M = 1$, $\lambda_0 = 60$, $\lambda_1 = 5$, $\mu = 1$, and $\lambda(u) = \lambda_0 - \lambda_1 u$.

be different, however, as far as their demand functions are concerned. It is not hard to see that a revenue-maximizing provider will generally charge different prices to the different classes, even in the case of unlimited resources ($R = \infty$). Prices will tend to be larger for those classes that have relatively inelastic demand, which is a well-known characteristic of monopolistic pricing.

Even though different classes are charged different prices, the optimal prices can be determined by solving a DP problem with one-dimensional (1-D) state space. This is because, once a call is admitted, its future behavior is independent of its class (because of μ and r being the same), and the state of the system is simply the number $n = n_1 + \dots + n_M$ of active calls. In particular, the relative reward function $h(\mathbf{N})$ is of the form $h(n)$, and can be easily computed by solving Bellman's equation. Once the function h has been computed, the optimal price for class i , when the state is $n < R$, can be found by maximizing

$$\lambda_i(u_i)u_i + \frac{\lambda_i(u_i)}{\nu}(h(n+1) - h(n)).$$

When it comes to welfare maximization, it turns out that the same price is charged to all classes. To see this, we first note that $h(n)$ depends only on the scalar $n = n_1 + \dots + n_M$, by our earlier argument. The optimal price for class i , at state $n < R$, is determined by maximizing

$$\lambda_i(u_i) \mathbf{E}[U_i | U_i \geq u_i] + \frac{\lambda_i(u_i)}{\nu}(h(n+1) - h(n)). \quad (12)$$

If $f_i(\cdot)$ is the probability density function of U_i , then (3) and (4) show that we need to maximize

$$\int_{u_i}^{u_i, \max} f_i(v)v dv + \frac{h(n+1) - h(n)}{\nu} \int_{u_i}^{u_i, \max} f_i(v) dv$$

with respect to u_i . Assuming an interior solution, the first order conditions yield

$$-u_i f_i(u_i) - \frac{f_i(u_i)}{\nu}(h(n+1) - h(n)) = 0$$

which leads to

$$u_i^*(n) = \frac{h(n) - h(n+1)}{\nu}, \quad \forall i.$$

V. STATIC PRICING STRATEGY

We say that a pricing policy is *static* if a fixed price vector \mathbf{u} is always in effect, independent of the state of the system. (Note

that this definition deviates somewhat from our earlier conventions, because a static price results in a constant arrival rate even if the required resources are unavailable, in which case calls are blocked.) Static prices are of interest because: (1) the computation of optimal dynamic prices increases exponentially with the number of classes and the capacity and (2) dynamic prices can be unattractive to users who may prefer facing a predictable, fixed, pricing structure.

Under a static pricing policy \mathbf{u} , the system evolves as a continuous-time Markov chain whose steady-state probability can be found in closed form [18]. The corresponding average revenue, denoted by $J(\mathbf{u})$, is given by

$$J(\mathbf{u}) = \sum_{i=1}^M \lambda_i(u_i) u_i (1 - \mathbf{P}_{\text{loss}}^i[\mathbf{u}]) \quad (13)$$

where $\mathbf{P}_{\text{loss}}^i[\mathbf{u}]$ denotes the steady-state probability $\mathbf{P}[\mathbf{N}'\mathbf{r} + r_i > R]$ that a call of class i gets rejected. Clearly, the *optimal static revenue* J_s , defined by

$$J_s = \max_{\mathbf{u} \in \mathcal{U}} J(\mathbf{u}) \quad (14)$$

satisfies $J_s \leq J^*$. It is of interest to determine the gap between these two quantities; we return to this later.

The loss probabilities $\mathbf{P}_{\text{loss}}^i[\mathbf{u}]$ can be efficiently computed; see, e.g., Kaufman [25] for a method with $O(MR)$ complexity, or Mitra *et al.* [26] for fast approximations. By optimizing over \mathbf{u} , an optimal static price can be usually computed with moderate effort.

For the case of welfare maximization, the same discussion applies, with $\lambda_i(u_i)u_i$ replaced by $\lambda_i(u_i)\mathbf{E}[U_i | U_i \geq u_i]$. Furthermore, if all classes share the same r_i and μ_i , it can be verified that the optimal price is the same for every class.

VI. AN UPPER BOUND ON THE OPTIMAL PERFORMANCE

To assess the degree of suboptimality of static or approximate dynamic policies, especially for high-dimensional problems where J^* is unknown, we develop an upper bound on J^* . One such bound is J_∞ (cf. Theorem 2), but this can be far from tight if the resource limitations are significant.

We use the inverse demand functions $u_i(\lambda_i)$, and let $F_i(\lambda_i)$ stand for the instantaneous reward rate when the present class i arrival rate is λ_i . Thus, $F_i(\lambda_i) = \lambda_i u_i(\lambda_i)$ or $F_i(\lambda_i) = \lambda_i \mathbf{E}[U_i | U_i \geq u_i(\lambda_i)]$, for the case of revenue or welfare maximization, respectively. We assume that the functions F_i are concave. (This property is true for both cases of revenue or welfare maximization, when the demand function $\lambda_i(u_i)$ is linear.) Let J_{ub} be the optimal value of the following nonlinear programming problem:

$$\begin{aligned} & \text{maximize} && \sum_i F_i(\lambda_i) \\ & \text{subject to} && \lambda_i = \mu_i n_i \quad \forall i \\ & && \sum_i n_i r_i \leq R. \end{aligned} \quad (15)$$

Theorem 6: If the functions F_i are concave, then $J^* \leq J_{\text{ub}}$.

Proof: Consider an optimal dynamic pricing policy. Without loss of generality, we can assume that the price u_i becomes large enough and the arrival rate $\lambda_i(u_i)$ is equal to zero, whenever the state is such that a class i call cannot be admitted. In this proof, we view λ_i and n_i as random variables, and use $\mathbf{E}[\cdot]$, to indicate expectation with respect to the steady-state distribution under this particular policy. At any time, we have $\sum_i r_i n_i \leq R$, which implies that $\sum_i \mathbf{E}[n_i] r_i \leq R$. Furthermore, Little's law implies that $\mathbf{E}[\lambda_i] = \mu_i \mathbf{E}[n_i]$. This shows that the $\mathbf{E}[n_i]$ and $\mathbf{E}[\lambda_i]$, $i = 1, \dots, M$, form a feasible solution of the problem (15). Therefore, using the concavity of F and Jensen's inequality

$$J_{\text{ub}} \geq \sum_i F_i(\mathbf{E}[\lambda_i]) \geq \sum_i \mathbf{E}[F_i(\lambda_i)] = J^*$$

where the last equality used the optimality of the policy under consideration. ■

Note that J_{ub} is very easy to compute, especially under our concavity assumption. It is an optimistic upper bound because it implicitly assumes that if the arrival rates are held constant to some values λ_i that satisfy $\sum_i \lambda_i r_i / \mu_i \leq R$, then no calls are blocked. Despite that, the optimal solution of (15) can form the basis of an approximately optimal static strategy, as will be discussed later.

If the functions F_i are not concave, then J_{ub} is not, in general, a valid upper bound, but we may proceed as follows. Consider a new problem with reward rates $\bar{F}_i(\lambda_i)$, instead of $F_i(\lambda_i)$, where \bar{F}_i is the smallest concave function that satisfies $\bar{F}_i(\lambda_i) \geq F_i(\lambda_i)$ for all λ_i . Let \bar{J}^* be the optimal average reward and let \bar{J}_{ub} be the optimal value in the maximization (15), when each F_i is replaced by \bar{F}_i . Using Theorem 6, applied to the new problem, we have

$$J^* \leq \bar{J}^* \leq \bar{J}_{\text{ub}}.$$

VII. LIMITING REGIMES

In this section, we consider a number of limiting regimes and discuss the nature of the resulting optimal policies. In many cases, static policies are asymptotically optimal. The most interesting regime involves a system with a large number of small users, which is considered first.

A. Many Small Users

If the capacity R is large compared to the bandwidth of a typical call, we expect that the laws of large numbers will take over, eliminate statistical fluctuations, and allow us to carry out an essentially deterministic analysis. To capture a situation of this nature, we start with a base system, involving a finite capacity R and finite demand functions $\lambda_i(u_i)$. We then scale the system through a proportional increase of capacity and demand.

More specifically, let $c \geq 1$ be a scaling factor. The scaled system has resources R^c , with $R^c = cR$, and demand functions $\lambda_i^c(u_i)$, given by $\lambda_i^c(u_i) = c\lambda_i(u_i)$. Note that the other parameters r_i and μ_i are held fixed. We will use a superscript c to denote various quantities of interest for the scaled system. It is easily seen that the optimal performance $J^{*,c}$, as well as the various bounds, such as J_{ub}^c , will increase roughly linearly with c ,

and for this reason, a meaningful comparison should first divide such quantities by c , as in the result that follows.

Theorem 7: Consider either the revenue or the welfare maximization problem, and assume that the functions $F_i(\lambda_i)$ (defined as in the preceding section) are concave. Then

$$\lim_{c \rightarrow \infty} \frac{1}{c} J_s^c = \lim_{c \rightarrow \infty} \frac{1}{c} J^{*,c} = \lim_{c \rightarrow \infty} \frac{1}{c} J_{\text{ub}}^c.$$

Proof: For simplicity, we only work with the revenue maximization formulation. The upper bound J_{ub}^c is obtained by maximizing $\sum_i c \lambda_i(u_i) u_i$, subject to the constraint

$$\sum_i \frac{c \lambda_i(u_i) r_i}{\mu_i} \leq cR. \quad (16)$$

It is easily seen that there exists an optimal solution $\mathbf{u}_{\text{ub}}^* = (u_{\text{ub},1}^*, \dots, u_{\text{ub},M}^*)$, which is independent of c , and that $J_{\text{ub}}^c = c J_{\text{ub}}^1$.

Fix some $\epsilon > 0$ and let us consider new static prices u_i^ϵ given by $u_i^\epsilon = u_{\text{ub},i}^* + \epsilon$. Let $J^c(\mathbf{u}^\epsilon)$ be the resulting average revenue. For every i such that $\lambda_i(u_{\text{ub},i}^*) > 0$, we have $\lambda_i(u_i^\epsilon) < \lambda_i(u_{\text{ub},i}^*)$. Let n_i^ϵ (respectively, $n_{i,\infty}^\epsilon$) be the random variable which is equal to the number of active class i calls, in steady-state, in the scaled system, under the prices u_i^ϵ , with capacity cR (respectively, with infinite capacity). Using that the $u_{\text{ub},i}^*$ satisfy (16), we obtain

$$\begin{aligned} & \mathbf{P} \left[\sum_i r_i n_{i,\infty}^\epsilon > cR - r_{\max} \right] \\ & \leq \mathbf{P} \left[\sum_i r_i n_{i,\infty}^\epsilon > \sum_i \frac{c \lambda_i(u_{\text{ub},i}^*) r_i}{\mu_i} - r_{\max} \right] \\ & = \mathbf{P} \left[\sum_i r_i \frac{n_{i,\infty}^\epsilon}{c} > \sum_i \frac{\lambda_i(u_{\text{ub},i}^*) r_i}{\mu_i} - \frac{r_{\max}}{c} \right] \end{aligned} \quad (17)$$

where $r_{\max} = \max_i r_i$.

Next note that $n_{i,\infty}^\epsilon$ is equal to the number of customers in an M/M/ ∞ queue with arrival rate $c \lambda_i(u_i^\epsilon)$ and service rate μ_i . Thus, $n_{i,\infty}^\epsilon$ is a Poisson random variable with parameter $c \lambda_i(u_i^\epsilon) / \mu_i$. As $c \rightarrow \infty$, the mean of $n_{i,\infty}^\epsilon / c$ converges to $\lambda_i(u_i^\epsilon) / \mu_i$ and the variance to zero. Consequently, it can be seen that the random variable $n_{i,\infty}^\epsilon / c$ converges in probability to the deterministic value $\lambda_i(u_i^\epsilon) / \mu_i$, which is less than $\lambda_i(u_{\text{ub},i}^*) / \mu_i$. It follows that the probability in the right-hand side of (17) converges to zero. We next compare n_i^ϵ and $n_{i,\infty}^\epsilon$. Comparing the number of customers in the two corresponding systems (one with capacity cR and the other with infinite capacity), and by defining the arrival processes on a common probability space, we conclude that for all sample paths n_i^ϵ is smaller than $n_{i,\infty}^\epsilon$. Hence

$$\begin{aligned} \mathbf{P}_{\text{loss}}^i[\mathbf{u}^\epsilon] &= \mathbf{P} \left[\sum_j r_j n_j^\epsilon > cR - r_i \right] \\ &\leq \mathbf{P} \left[\sum_j r_j n_j^\epsilon > cR - r_{\max} \right] \end{aligned}$$

$$\leq \mathbf{P} \left[\sum_j r_j n_{j,\infty}^\epsilon > cR - r_{\max} \right]$$

and the loss probabilities $\mathbf{P}_{\text{loss}}^i[\mathbf{u}^\epsilon]$ also converge to zero. Using (13) and (14), it follows that

$$\lim_{c \rightarrow \infty} \frac{1}{c} J_s^c \geq \lim_{c \rightarrow \infty} \frac{1}{c} J^c(\mathbf{u}^\epsilon) = \sum_i u_i^\epsilon \lambda_i(u_i^\epsilon).$$

This is true for any positive ϵ . We now let ϵ go to zero, in which case u_i^ϵ tends to $u_{\text{ub},i}^*$. Using the continuity of the demand function, we obtain

$$\lim_{c \rightarrow \infty} \frac{1}{c} J_s^c \geq \sum_i u_{\text{ub},i}^* \lambda_i(u_{\text{ub},i}^*) = J_{\text{ub}}^1.$$

On the other hand, $J_s^c \leq J^{*,c} \leq J_{\text{ub}}^c = c J_{\text{ub}}^1$, and the result follows. \blacksquare

B. Some Qualitative Conclusions for the Many Small Users Case

Theorem 7 and its proof indicate that when the system is large (many small users), and under the concavity assumption, approximate optimality is obtained by slightly modifying the static prices derived through the optimization in the definition of J_{ub} [cf. (15)].¹ It is thus of interest to study the nature of these static prices. The insights to follow can be very valuable in narrowing down the design space when considering more sophisticated approximations or adaptive methods.

1) Revenue Maximization: For the purposes of this discussion, it is more convenient to view the problem (15) as one involving optimization with respect to u_i , rather than λ_i . We write the resource constraint in the form $\sum_i \lambda_i(u_i) r_i / \mu_i \leq R$, and associate it with a nonnegative Lagrange multiplier q . Then, the class i price u_i is determined by maximizing $u_i \lambda_i(u_i) - q(\lambda_i(u_i) r_i) / \mu_i$. Assuming an interior solution, we differentiate and set the derivative to zero, to see that the maximizing prices satisfy

$$u_i = - \frac{\lambda_i(u_i)}{d\lambda_i(u_i)/du_i} + q \frac{r_i}{\mu_i}, \quad \forall i. \quad (18)$$

Thus, u_i consists of two parts. The first is the reciprocal of the demand elasticity, and illustrates how a monopolist will tend to increase revenue by overcharging inelastic demand. The second term is a usage-based charge. The quantity r_i / μ_i is the ‘‘volume’’ (product of bandwidth and time) consumed by a class i call. It is important to notice that this usage-based charge is determined by a single parameter q , common to all classes. Equally important, such a charge can be implemented without knowing μ_i , by setting a connection-time charge rate equal to qr_i .

In the special case of linear demand functions of the form $\lambda_i(u_i) = \lambda_{0,i} - \lambda_{1,i} u_i$, we can solve for u_i and obtain

$$u_i = \frac{\lambda_{0,i}}{2\lambda_{1,i}} + \frac{q r_i}{2\mu_i}. \quad (19)$$

¹Actually, this small modification is not necessary, but then a more complicated proof would be required.

2) *Welfare Maximization:* We proceed similar to the preceding case. After we introduce the Lagrange multiplier q , we need to maximize, with respect to u_i , the expression [cf. (4)]

$$\lambda_{0,i} \int_{u_i}^{u_{i,\max}} v f_i(v) dv - q \frac{\lambda_i(u_i) r_i}{\mu_i}.$$

By differentiating, we obtain a condition for optimality of an interior solution, of the form

$$\lambda_{0,i} u_i f_i(u_i) = -q \frac{r_i}{\mu_i} \frac{d\lambda_i(u_i)}{du_i}.$$

Using (3), this leaves us with

$$u_i = q \frac{r_i}{\mu_i}. \quad (20)$$

Thus, there is a single price for volume, and all classes are charged in proportion to the volume that they consume. (Such a property is also present in [3] although the setting is different, as it relates to elastic traffic.)

C. The Non-Concave Case

As discussed at the end of Section VI, J_{ub} is not, in general, a valid upper bound, in the absence of concavity. Rather, we first need to “convexify” the problem, by replacing the functions F_i by their concave counterparts \bar{F}_i , leading to the upper bound \bar{J}_{ub} . Unlike the concave case, the gap between J_s and \bar{J}_{ub} does not vanish in the limit of many small users, and static policies are no more asymptotically optimal. It turns out that asymptotically optimal policies, whose performance is close to \bar{J}_{ub} , are again easy to obtain, but their form is less appealing. Rather than developing a general theory, we focus on the single class case, and assume, without loss of generality that $r = 1$. The generalization to multiple classes is straightforward.

In the single class case, we have

$$J_{\text{ub}} = \max_{\lambda \leq \mu R} F(\lambda), \quad \bar{J}_{\text{ub}} = \max_{\lambda \leq \mu R} \bar{F}(\lambda).$$

If $\bar{J}_{\text{ub}} = J_{\text{ub}}$, the approach of Section VII–A provides a static policy whose performance approaches J_{ub} and we are done. Suppose now that $J_{\text{ub}} < \bar{J}_{\text{ub}}$. Let λ^* be an arrival rate that attains the maximum in the definition of \bar{J}_{ub} . Because of the way that \bar{F} has been defined (“convexification” of F), there exist λ_1 and λ_2 and some $\alpha \in (0, 1)$, such that

$$\lambda^* = \alpha \lambda_1 + (1 - \alpha) \lambda_2, \\ \bar{F}(\lambda^*) = \alpha F(\lambda_1) + (1 - \alpha) F(\lambda_2).$$

Let u_1, u_2 be prices such that $\lambda_1 = \lambda(u_1)$ and $\lambda_2 = \lambda(u_2)$, where $\lambda(\cdot)$ is the demand function. Without loss of generality, assume that $u_1 < u_2$, which leads to $\lambda_1 > \lambda_2$. Furthermore, it is not hard to show that in this case, we will have $\lambda_2 < \mu R < \lambda_1$. Indeed, if $\lambda_2 < \lambda_1 \leq \mu R$, then we have $\bar{J}_{\text{ub}} = \bar{F}(\lambda^*) \leq \max\{F(\lambda_1), F(\lambda_2)\} \leq J_{\text{ub}}$, which leads to a contradiction.

Consider a state-independent but time-varying pricing policy that sets the price to u_1 for a fraction α of the time and to u_2 for a fraction $1 - \alpha$ of the time, and switches between these two prices with very high frequency. In the limit, as the frequency

of switching increases, the system is faced with an arrival rate of λ^* and an expected reward per unit time (in the absence of blocking) equal to $\bar{F}(\lambda^*)$. In the many small user regime, the probability of blocking can be made arbitrarily small, through a small modification of the prices u_1 and u_2 , as in the proof of Theorem 7. Thus, the average performance of the system approaches $\bar{F}(\lambda^*) = \bar{J}_{\text{ub}}$, and we have asymptotic optimality.

An alternative implementation is to segment the user population into two classes, with a fraction α being assigned to the first class. A price (approximately equal to) u_i is then quoted to class i . What may be surprising with this scheme is that market segmentation is beneficial even if the two market segments have identical technical characteristics (μ and r), the *same demand elasticities*, and even if the objective is to maximize social welfare!

The schemes in the preceding two paragraphs are not realistic. If the price switches frequently between u_1 and u_2 , users will naturally try to “time the market” and place calls when the lower price is in effect, invalidating the Poisson arrival model. If on the other hand the market is segmented, users who have been assigned to the higher priced segment will either try to switch to the lower priced segment or to another provider with “fairer” practices.

Of course, we know that there exists an optimal dynamic pricing policy which is stationary (at all times, it is the same function of the state of the system), and which does not explicitly segment the market. Without carrying out a detailed analysis, it is not hard to guess the form of a close-to-optimal dynamic policy with these properties, in the many small user regime. We choose a threshold state n^* near R . We then set $u(n) \approx u_1$ for $n \leq n^*$, and $u(n) \approx u_2$ for $n > n^*$. It was noted earlier that $\lambda_2 < \mu R < \lambda_1$, and let us assume that n^* is close enough to R so that $\lambda_2 < \mu n^* < \lambda_1$. Then, the system will spend most of the time in the vicinity of n^* and the price will be switching very frequently between u_1 and u_2 , thus achieving the same effect as in the two schemes that were discussed earlier, including undesirable fluctuations at a fast time scale.

It is unclear whether nonconcave reward rates $F(\lambda)$ can arise in practice. If they do, the discussion above implies that near-optimality may not be practically attainable.

D. Light or Heavy Traffic

For simplicity, we only discuss revenue maximization and a single class. We use a demand function $c\lambda(u)$, and we take the limit as $c \rightarrow 0$ (light traffic), or $c \rightarrow \infty$ (heavy traffic). In light traffic, the resource constraint becomes immaterial, and the best static price for an unconstrained system, u_∞ , becomes asymptotically optimal. In heavy traffic, the system can be fully utilized as long as the price is slightly less than u_{\max} . We summarize the consequences of these observations in the theorem below. The proof is straightforward and is omitted.

Theorem 8: a) (**Light Traffic**) We have $\lim_{c \rightarrow 0} (J_s/c) = \lim_{c \rightarrow 0} (J^*/c) = \lim_{c \rightarrow 0} (J_\infty/c) = \lambda(u_\infty) u_\infty$.

b) (**Heavy Traffic**) We have $J^* \leq K \mu u_{\max}$, where $K = \lceil R/r \rceil$ and the inequality is asymptotically tight as $c \rightarrow \infty$. Moreover, for all ϵ such that $0 < \epsilon < u_{\max}$, the static price $u_{\max} - \epsilon$ achieves revenue J_ϵ that satisfies $\lim_{c \rightarrow \infty} J_\epsilon = (u_{\max} - \epsilon) K \mu$.

VIII. APPROXIMATELY OPTIMAL STATIC POLICIES

Computing exactly optimal static policies can be difficult when the number of classes is large, because of numerical problems with the computation of the loss probabilities and also because we need to carry out an optimization over a high-dimensional set of possible price vectors \mathbf{u} . We discuss a number of approaches to get around this difficulty.

- 1) We can use the static prices suggested by the optimization problem in (15). These will be of the form (18) or (20), where q is the optimal Lagrange multiplier associated with the resource constraint in (15).
- 2) More generally, we can use prices of the form (18) or (20), and tune q to optimize performance. This option may be called for if we are not quite in the limiting regime and the multiplier obtained from the maximization (15) cannot be fully trusted. For example, we can simulate the system for different choices of q and pick the best value. Alternatively, we can estimate the gradient of the performance metric with respect to q , and iteratively adjust q in the course of a single simulation, using for example, the methods of Marbach and Tsitsiklis [27].
- 3) We can add still more flexibility by viewing each u_i as an independent tunable parameter (as opposed to using a single parameter q), and then either use fast approximations of the blocking probabilities (as in [26]), or employ a simulation-based method, as discussed above.

An advantage of the simulation-based methods described in items 2 and 3 is that they may be carried on-line. One can then explore model-free versions that do not assume explicit knowledge or estimation of demand. This would lead to methods that can adapt the static prices and track slow changes in the demand functions.

IX. APPROXIMATELY OPTIMAL DYNAMIC POLICIES

The computation of optimal dynamic policies becomes intractable even with a moderate number of classes. If we are far from the asymptotic regime where static prices are near-optimal, it becomes of interest to explore methods that lead to approximately optimal dynamic policies. Possible approaches include the following.

- 1) Construct a low-dimensional Markov model that approximates the exact model, and use an optimal policy for the approximate problem. For example, Subramanian *et al.* [14] have used a 1-D approximate Markov model to address an airline yield management problem. In our case, we can use a two-dimensional state (x, y) with components $x = \sum_i n_i r_i$ and $y = \sum_i n_i r_i / \mu_i$. Note that y measures the currently occupied volume, and has been shown to be an important quantity by our analysis in Section VII-B. It is not clear whether a 1-D model could capture both the resource constraints and the importance of volume.
- 2) Work with the original model, but use a parametric representation of the reward function, e.g., of the form

$$\tilde{h}(\mathbf{N}, \boldsymbol{\theta}) = \theta_0 + \sum_i \theta_i n_i + \sum_{ij} \theta_{ij} n_i n_j. \quad (21)$$

One can employ any one of many available approximate DP methods (“neuro-dynamic programming” or “reinforcement learning” [28], [29]), all of which aim at setting the values of the parameters $\boldsymbol{\theta}$ in a way that results in good performance. This methodology has been used successfully in a wide range of resource allocation problems, including some complex admission control and routing problems [30]. In Section X-C, we apply such an approach to derive an approximately optimal dynamic policy for large-scale problems. We use the parametric representation (21) and apply a method originating in [31] and described in [28].² Define $(Q\mathbf{h})(\mathbf{N}, \mathbf{u})$ to be equal to the expression in the right-hand side of the Bellman equation (7), when prices are fixed to \mathbf{u} . It is well known [23] that a solution to (7) can be obtained by solving the linear programming problem:

$$\begin{aligned} & \text{minimize } J \\ & \text{subject to } J + h(\mathbf{N}) \geq (Q\mathbf{h})(\mathbf{N}, \mathbf{u}) \quad \forall \mathbf{N}, \quad \forall \mathbf{u} \in \mathcal{U}_d \end{aligned} \quad (22)$$

where we discretize prices and \mathcal{U}_d is a discrete subset of \mathcal{U} . Using the parametric representation (21) of \mathbf{h} , we obtain a linear programming problem with $O(M^2)$ decision variables, but with a very large number of constraints. To address this, we apply a cutting-plane method guided by a simulation of the system. We initialize the method by solving (22) with a small number of constraints corresponding to values (\mathbf{N}, \mathbf{u}) uniformly drawn from the state space and \mathcal{U}_d . The resulting $(J, \boldsymbol{\theta})$ leads to a policy through the maximization of the right-hand side of Bellman’s equation. We simulate this policy and for each state \mathbf{N} encountered in the simulation, we check whether the constraints

$$J + \tilde{h}(\mathbf{N}) \geq (Q\tilde{\mathbf{h}})(\mathbf{N}, \mathbf{u})$$

are satisfied for all $\mathbf{u} \in \mathcal{U}_d$. We solve again the linear programming problem, with all violated constraints at state \mathbf{N} included, to obtain a new policy, and continue the simulation from \mathbf{N} with the new policy. Thus, at each iteration we “improve” $\tilde{\mathbf{h}}$ by adding constraints violated by the previous optimal solution.

- 3) Consider dynamic policies that depend on a small number of parameters and tune those parameters to optimize performance. The main idea here is the same as in items 2 and 3 of the preceding section. For example, guided by the insights of Section VII-B, we could use prices of the form $u_i(\mathbf{N}) = (r_i / \mu_i)(\theta_0 + \theta_1 b + \theta_2 b^2)$, where $b = \sum_i r_i n_i$.

X. NUMERICAL COMPARISONS BETWEEN STATIC AND DYNAMIC POLICIES

In this section, we numerically compare the performance of the optimal dynamic and static pricing policies. We will see that the static policy offers its substantial implementation advantage at a modest performance cost, in agreement with the theoretical

²Other approximate DP methods are also possible, but we took this as an opportunity to test a method for which there is practically no experience.




Fig. 2. For the example of Fig. 1, we plot the fraction of time $P[n]$ that the system is in state n , in steady-state, under the dynamic and the static pricing policies, respectively.

conclusions from Section VII. We only consider revenue maximization problems. The results for welfare maximization would not be much different.

A. Single Class Example

Among the objectives of pricing policies is to provide incentives so that calls modify their arrival patterns and reduce demand during congested periods, leading to more efficient resource utilization. To assess the extent to which this goal is achieved, we depict in Fig. 2, and for the example of Fig. 1, the steady-state system occupancy under the dynamic and the static pricing policies, respectively. We observe that the curve for the static policy closely approximates the one for the dynamic policy. Under both policies the system spends most of the time with about 25 customers present (i.e., about 83% utilization). As expected, dynamic pricing leads to better rationing of the resources, with more “impulsive” system occupancy.

B. Two-Class Example

We next turn our attention to a system with two service classes with linear demand functions of the form $\lambda_i(u_i) = \lambda_{0,i} - u_i \lambda_{1,i}$. The objective is to demonstrate that the optimal static policy is not significantly inferior to the optimal dynamic one. Moreover, we want to illustrate that by just setting appropriate static prices, the provider can exploit class characteristics to maximize revenues. The results for the two-class system are reported in Table II. The input parameters for the calculations in Table II are given in Table I. These results suggest that the optimal static pricing is very close to optimal, the suboptimality gap being less than 2%. Note that in light traffic (top rows of Table II) the static policy performs very well. The same is true when the system is highly congested (bottom rows of Table II). This is to be expected in view of the results in Section VII-D. Moreover, J_{ub} exhibits the same behavior (i.e., it is especially tight in light and heavy traffic).

Consider, for example, the case in the first row of Table II. Note that class one customers require 4 times as much bandwidth and stay for twice as long, on the average. That is, class one customers are “fat and slow” while class two customers are “slim and fast”. Although both classes have comparable ρ 's, it turns out that under the optimal static policy class one suffers a loss probability of 3.6%, while the corresponding class two value is 0.79%. Moreover the optimal static prices are $(u_{s,1}, u_{s,2}) = (7.08, 5.24)$. Let us recall the theory of

TABLE I
INPUT PARAMETERS FOR THE COMPUTATIONS IN TABLE II.
WE USED $R = 155, r_1 = 4, r_2 = 1, \mu_1 = 1,$ AND $\mu_2 = 2$. HERE,
 $\rho_i = \lambda_{0,i} r_i / (R \mu_i)$ IS A MEASURE OF THE LOAD PRESENTED BY
EACH CLASS i

	$\lambda_{0,1}$	$\lambda_{1,1}$	ρ_1	$\lambda_{0,2}$	$\lambda_{1,2}$	ρ_2
Case 1	40	4	1.032	350	35	1.129
Case 2	40	4	1.032	500	50	1.613
Case 3	80	8	2.064	350	35	1.129
Case 4	80	8	2.064	500	50	1.613
Case 5	160	16	4.128	1280	128	4.129
Case 6	320	32	8.256	2560	256	8.258
Case 7	640	64	16.512	5120	512	16.516

Section VII-B and, in particular, (19), which suggests prices of the form $(u_1, u_2) = (5 + 2q, 5 + (q/4))$. By setting $q = 1.04$, we see that these latter prices are in close agreement with the optimal static prices. That is, the upper bound calculation can capture the form of the optimal static prices. The differences between $u_{ub,i}$ and $u_{s,i}$ observed in Table II can be attributed to the fact that the upper bound computation does not provide an accurate estimate of q .

It turns out that class two contributes 91.6% to the total revenue. It is evident that class one needs more resources than class two, and the system can focus on the second class and make most of its revenue from it. Indeed, the static price $u_{s,2}$ for class two is fairly close to the optimal infinite capacity price $u_{\infty,2}$. The steady-state system occupancy under the optimal static policy is depicted in Fig. 3.

C. Large-Scale Examples: Approximate Dynamic Policies

In this subsection, we turn our attention to relatively large-scale problems. The computation of optimal static and dynamic prices becomes computationally prohibitive as the state space grows, thus, we will resort to the approximation methods outlined in Sections VIII and IX.

In Table IV, we report approximate DP results for a number of problems including large-scale ones. We have used the approach outlined in item 2 of Section IX. Table III reports the corresponding input parameters. We observe that the performance of the approximate dynamic policy is fairly close to the optimal even for large scale problems (since it is close to the upper bound). The form of the approximate dynamic policy for Case 3 is depicted in Fig. 4. Notice that the resulting prices have a “staircase” character, which is due to the discretization of the control space.

Although this particular approximate dynamic policy may underperform the optimal static pricing policy, there are some reasons that can make it attractive.

- 1) The static pricing policy treats all users of the same class equally, in the sense that they face the same probability of getting a busy signal. On the other hand, the approximate dynamic policy allows customers who place higher value on their connection than others to increase their chance of getting connected by paying a higher price.
- 2) Dynamic policies can be more robust to errors in the demand estimation. We elaborate further on this point in the next section.

TABLE II
 NUMERICAL RESULTS FOR THE TWO-CLASS SYSTEM. WE DENOTE BY J^* (RESP. J_s) THE OPTIMAL DYNAMIC (RESP. OPTIMAL STATIC) REVENUE. J_{ub} DENOTES THE OPTIMAL VALUE OF THE PROBLEM IN (15), $(u_{ub,1}, u_{ub,2})$ THE PRICES SUGGESTED BY THE UPPER BOUND COMPUTATION, AND $(u_{s,1}, u_{s,2})$ THE OPTIMAL STATIC PRICES

	J_s	J^*	J_{ub}	$\frac{J^* - J_s}{J^*} \times 100\%$	$u_{ub,1}$	$u_{s,1}$	$u_{ub,2}$	$u_{s,2}$
Case 1	945.79	952.63	972.85	0.72%	5.69	7.08	5.09	5.24
Case 2	1270.4	1281.65	1317.32	0.88%	7.62	8.74	5.33	5.42
Case 3	965.33	977.28	1012.43	1.22%	7.71	8.23	5.34	5.38
Case 4	1273.9	1288.97	1329.72	1.17%	8.7	9.26	5.46	5.48
Case 5	2206.1	2235.13	2349.22	1.30%	10	10	7.58	7.53
Case 6	2588.9	2613.36	2724.60	0.94%	10	10	8.79	8.64
Case 7	2804.1	2820.47	2912.30	0.58%	10	10	9.39	9.24




Fig. 3. Steady-state system occupancy under the optimal static policy for the two-class system.

TABLE III
 INPUT PARAMETERS FOR THE RESULTS OF TABLE IV WE CONSIDERED TWO-CLASS SYSTEMS WITH DEMAND FUNCTIONS OF THE FORM $\lambda_i(u_i) = \lambda_{0,i} - u_i \lambda_{1,i}$ AND $r_1 = 4, r_2 = 1, \mu_1 = 1, \mu_2 = 2$

	R	$\lambda_{0,1}$	$\lambda_{1,1}$	ρ_1	$\lambda_{0,2}$	$\lambda_{1,2}$	ρ_2
Case 1	10	40	4	16	350	35	17.5
Case 2	155	40	4	1.03	350	35	1.13
Case 3	155	70	4	1.81	550	35	1.78
Case 4	1550	400	40	1.03	3500	350	1.13
Case 5	8500	400	40	0.19	35000	3500	2.06

XI. TIME-OF-DAY AND ADAPTIVE PRICING

We have assumed throughout that the statistics of the arrival and service processes are stationary, which led to the development of stationary pricing policies (dynamic and static). In practice, these statistics typically vary with the time of the day. We expect, however, that they are slowly varying, and that one can sufficiently approximate $\lambda(\mathbf{u}, t)$ and $\mu(t)$ by piecewise constant functions of time. In particular, one can define a number of L time intervals s_1, \dots, s_L that span a 24-h period (e.g., $L = 3, s_1 = [9:00, 17:00], s_2 = [17:00, 1:00], s_3 = [1:00, 9:00]$) such that for all $k = 1, \dots, L$

$$\lambda(\mathbf{u}, t) \approx \lambda^k(\mathbf{u}), \quad \mu(t) \approx \mu^k, \quad t \in s_k, \quad (23)$$

where $\lambda^k(\mathbf{u})$ and μ^k are constant functions of time. Thus, to implement a static pricing policy, it suffices to calculate the static price for each such interval (as a function of $\lambda^k(\mathbf{u})$ and μ^k). The resulting policy is a *time-of-day pricing policy* with L prices. Typical demand patterns in the Internet suggest that a relatively small value of L (e.g., 3 or 4) can yield a good approximation of traffic statistics.

Alternatively, it is possible that demand undergoes slow but unpredictable changes. In that case, we can let the static prices change in an adaptive manner, e.g., by tuning the u_i 's (or the single parameter q of Section VII-B) according to a stochastic iterative method. Without such adaptation, and if an incorrect demand model is used, the resulting prices can be far from optimal.

It is interesting to note that a significant degree of adaptivity can also be accomplished using dynamic (state-dependent) but stationary (not explicitly depending on time) pricing. For example, consider the single-class case under the objective of welfare maximization, with $r = 1$, and let $u(n) = 0$ for $n < 9R/10$. For n between $9R/10$ and R , let $u(n)$ increase smoothly from 0 to u_{max} . Then, it is easy to see that in the many-small-users regime of Section VII-B, the utilization of the system will be at least 0.9, and we will have near optimality even if we do not know the true demand function. This argument indicates that even though dynamic pricing does not perform much better than static pricing when an exact model is available, it provides a degree of adaptivity when a demand model is unavailable.

XII. CONCLUSIONS AND EXTENSIONS

We have introduced and studied a model for optimal congestion-dependent (dynamic) pricing of network services, with the twofold objective of developing approximately optimal methods as well as useful insights. We have carried out the analysis in both a revenue maximization and a welfare maximization setting. We explored a number of alternatives such as the computation of the exact optimum and several approximations, and have provided a comparison with congestion-independent (static) pricing. Some of the most important conclusions are that static pricing can come very close to optimality and that a single price parameter q (volume charge) may suffice, especially when typical calls are relatively small. This leads to the familiar time-of-day pricing policies. We also saw that a revenue-maximizing provider may set substantially different prices for two services even if they have very similar resource requirements. This is consistent to what is happening in other industries (e.g., in air travel all passengers receive essentially the same service but can pay very different prices). However, we established that this is usually not the case when the objective is to maximize social welfare.

While we have considered a single shared resource, a similar model is possible involving several shared resources, each

TABLE IV
APPROXIMATE DP RESULTS FOR THE CASES OF TABLE III. \bar{J} DENOTES THE REVENUE GENERATED BY THE APPROXIMATE POLICY. FOR CASES 4 AND 5 IT IS COMPUTATIONALLY INTRACTABLE TO OBTAIN THE OPTIMAL STATIC AND DYNAMIC POLICIES

	\bar{J}	J_s	J^*	J_{ub}	$\frac{J_{ub}-\bar{J}}{J_{ub}} \times 100\%$	$u_{ub,1}$	$u_{s,1}$	$u_{ub,2}$	$u_{s,2}$
Case 1	159.54	163.73	164.63	188.57	15.39%	10.0	10	9.43	8.9
Case 2	920.53	945.79	952.63	972.85	5.38%	5.69	7.08	5.09	5.24
Case 3	2074.44	2164.4	2189.2	2260.87	8.25%	15.49	16.55	8.7	8.73
Case 4	8956.29	-	-	9728.5	7.94%	5.69	-	5.09	-
Case 5	85430.68	-	-	87772.28	2.67%	7.77	-	5.35	-




Fig. 4. Approximate dynamic prices for Case 3 of Table IV. Graphs (a) and (b) depict prices for classes one and two, respectively.

one resulting in a separate linear constraint on the state space. The case where service is delivered by a network, with each call using one or more links, falls in this category. Many other situations (see also [32]) can fit into this kind of model. The methods described in this paper can all be extended, in principle, although the increased complexity may require further approximations, e.g., requiring that the price of a call is equal to the sum of the prevailing prices for the different resources that it consumes.

APPENDIX

We provide here an example that shows that optimal prices can decrease with \mathbf{N} . Consider a system with two service classes ($M = 2$), in which R and $\mathbf{r} = (r_1, r_2)$ are such that the state space only contains the states $(0, 0), (0, 1), (1, 0)$, and $(2, 0)$. For an example, consider the case where $R = 10$ and $\mathbf{r} = (5, 8)$.

Consider the following demand functions: $\lambda_1(u_1) = 1 - u_1, u_1 \in [0, 1]$, and $\lambda_2(u_2) = 1 - (u_2/2), u_2 \in [0, 2]$. Let $\mu_i = \mu, i = 1, 2$, where μ is a large constant, much larger than 1. Note that the rate ν in the uniformized chain can be taken equal to $2\mu + 2$.

Because μ is much larger than the arrival rate, the system spends most of the time at state $(0, 0)$. Furthermore, if the system is started instead at some nonzero state \mathbf{N} , the resulting

expected total loss in revenue $(h(0, 0) - h(\mathbf{N}))/\nu$ is of the order of $1/\mu$. (This is proportional to the expected number of lost calls before the state becomes $(0, 0)$.) If we set $h(0, 0) = 0$, we then see that $h(\mathbf{N}) = O(1)$ for every other state. Using this fact, and writing down Bellman's equation for state $(0, 0)$, we see that the terms involving $h(\cdot)$ can be neglected (because they are multiplied by the factor $1/\nu$ and $h(\mathbf{N})/\nu = O(1/\mu)$). We obtain $u_1^*(0, 0) \approx 1/2, u_2^*(0, 0) \approx 1, J^* \approx 3/4$. Approximate equality here means that we are ignoring terms of the order of $1/\mu$.

Bellman's equation at state $\mathbf{N} = (1, 0)$ is of the form

$$\frac{3}{4} \approx \max_{u_1} \left[u_1(1 - u_1) + \frac{1 - u_1}{\nu} (h(2, 0) - h(1, 0)) + \frac{\mu}{\nu} (h(0, 0) - h(1, 0)) \right].$$

Since $\nu = O(\mu), \mu/\nu \approx 1/2$, and $h(0, 0) = 0$, we obtain $h(1, 0) \approx -1$.

Similarly, Bellman's equation at state $\mathbf{N} = (2, 0)$ is of the form

$$\frac{3}{4} \approx \frac{2\mu}{\nu} (h(1, 0) - h(2, 0))$$

and $h(2, 0) \approx -7/4$.

The above calculations show the lack of concavity of h as a function of n_1 , since $h(1, 0) - h(0, 0) \approx -1 < h(2, 0) - h(1, 0) \approx -3/4$. The optimal class 1 price at state $(n_1, 0)$, for $n_1 = 0, 1$, is determined by maximizing

$$u_1(1 - u_1) + \frac{1 - u_1}{\nu} (h(n_1 + 1, 0) - h(n_1, 0)).$$

In our example, $h(n_1 + 1, 0) - h(n_1, 0)$ increases with n_1 , and it is easily seen that the optimal u_1 decreases with n_1 .

The intuition behind this counterexample is the following. A single class one arrival blocks class two calls over a period of expected length $1/\mu$. Two class one arrivals block class two calls over a period of expected length $3/(2\mu)$, and the corresponding loss is less than twice the loss caused by a single arrival.

ACKNOWLEDGMENT

The authors would like to thank M. Caramanis, B. Prabhakar, and the participants of the Finland Workshop (June 1998) of the Internet Telephony Consortium for some useful discussions, as well as H. Huang and Y. Liu for helping with some of the numerical results. The second author is grateful to C. Courcoubetis and G. Stamoulis for hosting an extended visit to the Institute for Computer Science, Crete, Greece, where this work was partly carried out, as well as for several discussions.

REFERENCES

- [1] B. C. Smith, J. F. Leimkuhler, and R. M. Darrow, "Yield management at American Airlines," *Interfaces*, vol. 22, no. 1, pp. 8–31, 1992.
- [2] J. K. MacKie-Mason and H. R. Varian, "Pricing the Internet," in *Public Access to the Internet*, B. Kahin and J. Keller, Eds. Englewood Cliffs, NJ: Prentice-Hall, 1994.
- [3] F. P. Kelly, "Charging and rate control for elastic traffic," *Eur. Trans. Telecommun.*, vol. 8, pp. 33–37, 1997.
- [4] F. P. Kelly, A. K. Maulloo, and D. K. H. Tan, "Rate control for communication networks: Shadow prices, proportional fairness and stability," *J. Oper. Res. Soc.*, vol. 49, pp. 237–252, 1998.
- [5] R. J. Gibbens and F. P. Kelly, "Resource pricing and the evolution of congestion control," *Automatica*, vol. 35, no. 12, pp. 1969–1985, 1999.
- [6] D. D. Clark, "An Internet model for cost allocation and pricing," in *Internet Economics*, L. McKnight and J. Bailey, Eds. Cambridge, MA: MIT Press, 1997.
- [7] F. P. Kelly, "Notes on effective bandwidths," in *Stochastic Networks: Theory and Applications*, S. Zachary, I. B. Ziedins, and F. P. Kelly, Eds. London, U.K.: Oxford Univ. Press, 1996, vol. 9, pp. 141–168.
- [8] D. Bertsimas, I. Ch. Paschalidis, and J. N. Tsitsiklis, "On the large deviations behavior of acyclic networks of G/G/1 queues," *Ann. Appl. Prob.*, vol. 8, no. 4, pp. 1027–1069, 1998.
- [9] —, "Asymptotic buffer overflow probabilities in multiclass multiplexers: An optimal control approach," *IEEE Trans. Automat. Contr.*, vol. 43, pp. 315–335, March 1998.
- [10] I. Ch. Paschalidis, "Class-specific quality of service guarantees in multimedia communication networks," *Automatica*, vol. 35, no. 12, pp. 1951–1968, 1999.
- [11] F. P. Kelly, "On tariffs, policing and admission control for multiservice networks," *Oper. Res. Lett.*, vol. 15, pp. 1–9, 1994.
- [12] C. Courcoubetis, F. P. Kelly, and R. Weber, Measurement-based charging in communications networks. preprint.
- [13] S. A. Lippman and S. Stidham, "Individual versus social optimization in exponential congestion systems," *Oper. Res.*, vol. 25, pp. 233–247, 1977.
- [14] J. Subramanian, S. Stidham, and C. J. Lautenbacher, "Airline yield management with overbooking, cancellations, and no-shows," 1997. preprint.
- [15] G. Gallego and G. van Ryzin, "A multiproduct dynamic pricing problem and its applications to network yield management," *Oper. Res.*, vol. 45, no. 1, pp. 24–41, 1997.
- [16] P. B. Key, "Optimal control and trunk reservation in loss networks," *Prob. Eng. Inform. Sci.*, vol. 4, pp. 203–242, 1990.
- [17] T. J. Ott and K. R. Krishnan, "Separable routing: A scheme for state-dependent routing of circuit switched telephone networks," *Ann. Oper. Res.*, vol. 35, pp. 43–68, 1992.
- [18] K. W. Ross, *Multiservice Loss Models for Broadband Telecommunication Networks*. New York: Springer-Verlag, 1995.
- [19] S. H. Low and P. P. Varaiya, "A new approach to service provisioning in ATM networks," *IEEE/ACM Trans. Networking*, vol. 1, pp. 547–553, Oct. 1993.
- [20] R. Edell and P. Varaiya, "Demand for quality-differentiated network services," in *Proc. 36th IEEE Conf. on Decision and Control*, San Diego, CA, Dec. 1997, pp. 1443–1448.
- [21] C. A. Courcoubetis and M. I. Reiman, "Pricing in a Large Single Link Loss System," in *Teletraffic Engineering in a Competitive World*, P. Key and D. Smith, Eds. Amsterdam, The Netherlands: Elsevier, 1999, pp. 737–746.
- [22] P. J. Schweitzer, "A Brouwer fixed-point mapping approach to communicating Markov decision processes," *J. Math. Anal. Applicat.*, vol. 123, pp. 117–130, 1987.
- [23] D. P. Bertsekas, *Dynamic Programming and Optimal Control*. Belmont, MA: Athena Scientific, 1995, vol. I and II.
- [24] K. W. Ross and D. Tsang, "Optimal circuit access policies in an ISDN environment," *IEEE Trans. Commun.*, vol. 37, pp. 934–940, 1989.
- [25] J. S. Kaufman, "Blocking in a shared resource environment," *IEEE Trans. Commun.*, vol. COM-29, pp. 1474–1481, 1981.
- [26] D. Mitra, J. A. Morrison, and K. G. Ramakrishnan, "ATM network design and optimization: A multirate loss network framework," *IEEE/ACM Trans. Networking*, vol. 4, pp. 531–543, Aug. 1996.
- [27] P. Marbach and J. N. Tsitsiklis, "Simulation-based optimization of Markov reward processes," Cambridge, MA, Tech. Rep. LIDS-P-2411, LIDS, MIT, Feb. 1998.
- [28] D. P. Bertsekas and J. N. Tsitsiklis, *Neuro-Dynamic Programming*. Belmont, MA: Athena Scientific, 1996.
- [29] R. Sutton and A. Barto, *Reinforcement Learning*. Cambridge, MA: MIT Press, 1998.
- [30] P. Marbach, O. Mihatsch, and J. N. Tsitsiklis, "Call admission control and routing in integrated service networks using reinforcement learning," presented at the 37th Conf. on Decision and Control, Tampa, FL, Dec. 1998.
- [31] P. J. Schweitzer and A. Seidman, "Generalized polynomial approximations in Markovian decision processes," *J. Math. Anal. Applicat.*, vol. 110, pp. 568–582, 1985.
- [32] F. P. Kelly, "The Clifford Paterson lecture, 1995: Modeling communication networks, present and future," *Phil. Trans. R. Soc. London A*, vol. 354, pp. 437–463, 1996.



Ioannis Ch. Paschalidis (M'96) was born in Athens, Greece, in 1968. He received the Diploma degree in electrical and computer engineering from the National Technical University of Athens, Athens, Greece, in 1991, and the S.M. and Ph.D. degrees in electrical engineering and computer science from the Massachusetts Institute of Technology (MIT), Cambridge, in 1993 and 1996, respectively.

He was a Post-Doctoral Associate at the Laboratory for Information and Decision Systems, MIT, during the summer of 1996. Since September 1996, he has been with Boston University, Boston, MA, where he is Associate Professor of manufacturing engineering. His research interests include the analysis and control of stochastic systems with main applications in manufacturing systems and communication networks.

Dr. Paschalidis has received the second prize in the 1997 George E. Nicholson paper competition by INFORMS, and has been elected full member of Sigma Xi (MIT Chapter). He has served on the program committees of the 1997 INFORMS Applied Probability Conference, the 1998 IEEE Conference on Decision and Control, the organizing committee of 1998 Information Theory Symposium, and is on the program committee of SPIE 2000 Voice, Video, and Data Communications Symposium.



John N. Tsitsiklis (M'83–SM'97–F'99) was born in Thessaloniki, Greece, in 1958. He received the B.S. degree in mathematics in 1980, and the B.S., M.S. and Ph.D. degrees in electrical engineering, all from the Massachusetts Institute of Technology (MIT), Cambridge, in 1980, 1981 and 1984, respectively.

He was an Acting Assistant Professor of electrical engineering at Stanford University, Stanford, CA, during the academic year 1983–1984. Since 1984, he has been with MIT, where he is currently Professor of Electrical Engineering and Computer Science. He has served as Acting Codirector of the MIT Laboratory for Information and Decision Systems (Spring 1996 and 1997). He has also been a visitor with the Department of Electrical Engineering and Computer Science at the University of California, Berkeley, and the Institute for Computer Science in Iraklion, Greece. His research interests are in the fields of systems, optimization, control, and operations research. He has written more than 80 journal papers in these areas. He is a coauthor (with D. Bertsekas) of *Parallel and Distributed Computation: Numerical Methods* (Englewood Cliffs, NJ: Prentice-Hall, 1989) and *Neuro-Dynamic Programming* (Belmont, MA: Athena Scientific, 1996), and *Introduction to Linear Optimization* with D. Bertsimas (Belmont, MA: Athena Scientific, 1997). He is an associate editor of *Applied Mathematics Letters* and a past editor of *Automatica*.

Dr. Tsitsiklis has been a recipient of an IBM Faculty Development Award (1983), an NSF Presidential Young Investigator Award (1986), an Outstanding Paper Award by the IEEE Control Systems Society, the M.I.T. Edgerton Faculty Achievement Award (1989), the Bodossakis Foundation Prize (1995), and the INFORMS/CSTS prize (1997). He was a plenary speaker at the 1992 IEEE Conference on Decision and Control. He has been an associate editor of the IEEE TRANSACTIONS ON AUTOMATIC CONTROL.